

ABSTRACT OF THE INVENTION

A method is described for protecting dirty data in cache memories in a cost-effective manner. When an instruction to write data to a memory location is received, and that memory location is being cached, the data is written to a plurality of cache lines, which are referred to as duplicate cache lines. When the data is written back to memory, one of the duplicate cache lines is read. If the cache line is not corrupt, it is written back to the appropriate memory location and marked available. In one embodiment, if more duplicate cache lines exist, they are invalidated. In another embodiment, the other corresponding cache lines may be read for the highest confidence of reliability, and then marked clean or invalid.